arXiv:1402.4145v3 [quant-ph] 23 Aug 2015

Unbounded entanglement in nonlocal games

Laura Mančinska* Thomas Vidick†

August 25, 2015

Abstract

Quantum entanglement is known to provide a strong advantage in many two-party distributed tasks. We investigate the question of how much entanglement is needed to reach optimal performance. For the first time we show that there exists a purely classical scenario for which no finite amount of entanglement suffices. To this end we introduce a simple two-party nonlocal game H, inspired by Lucien Hardy's paradox. In our game each player has only two possible questions and can provide bit strings of any finite length as answer. We exhibit a sequence of strategies which use entangled states in increasing dimension d and succeed with probability $1 - O(d^{-c})$ for some $c \geq 0.13$. On the other hand, we show that any strategy using an entangled state of local dimension d has success probability at most $1 - \Omega(d^{-2})$. In addition, we show that any strategy restricted to producing answers in a set of cardinality at most d has success probability at most $1 - \Omega(d^{-2})$. Finally, we generalize our construction to derive similar results starting from any game G with two questions per player and finite answers sets in which quantum strategies have an advantage.

1 Introduction

Entanglement plays a key role in quantum information processing. The almost unnaturally strong correlations it implies were initially seen as a weird, if not undesirable [EPR35], feature of quantum mechanics. Yet more recently entanglement is increasingly being thought of as a distributed resource that allows cooperating parties to accomplish otherwise impossible tasks, such as unconditionally secure cryptography [Eke91], randomness certification [Col06, PAM+10] and expansion [VV12, CY14], or classical communication with improved efficiency [BW92]. All these tasks are purely classical scenarios in which the use of shared entanglement provides a strong advantage. It is thus natural to ask how much of this new resource is needed in order to achieve optimal performance. The question arises in areas such as quantum Shannon theory [Wil13], communication complexity [BCMdW10], nonlocality [BCP+14], and many others. Perhaps surprisingly, very few general results are known. Specific examples have been constructed to establish lower bounds on the amount of entanglement required. Yet the problem of proving upper bounds is a recurrent sticking point and few such bounds are known; two-player XOR games provide a rare exception [CHTW04].

This unfortunate state of affairs has led to the question of whether such bounds exist in principle. The question can be abstractly formulated as follows:

*Centre for Quantum Technologies, National University of Singapore. Email: laura@locc.la
†Simons Institute, University of California at Berkeley. Email: vidick@csail.mit.edu

Is it possible to associate an integer $d(\mathcal{S}) \in \mathbb{Z}_+$ to any finite size $\mathcal{S} \in \mathbb{Z}_+$ so that any distributed task $\mathcal{T}$ of size $\mathcal{S}$ can be completed with vanishing error probability using entangled states of local dimension at most $d(\mathcal{S})$? ($\star$)

We may also consider a less demanding finite-precision variant of the above question. That is, we could require to achieve $\mathcal{T}$ *approximately*, within some finite precision ε, and allow $d = d(\mathcal{S}, \varepsilon)$ to depend on ε as well. (For now we are purposefully leaving terms such as complexity, precision, etc., loosely defined; they will be made more concrete in the coming paragraphs.)

We study both questions in the context of nonlocal games. In such games, two separated parties are provided with questions x and y respectively, chosen according to a pre-specified distribution $\pi(x, y)$. Without communicating they must respectively provide answers a and b. The task $\mathcal{T}$ is to maximize the probability that their answers satisfy a pre-determined criterion $V(a, b|x, y) = 1$. The entangled value, $\omega^*(G)$, of such a game $G = (V, \pi)$ is the largest success probability achievable using finite-dimensional entangled states (see Section 2 for precise definitions). It is known that for any $d \in \mathbb{Z}^+$ there exists a nonlocal game for which any strategy attaining the optimal success probability requires measurements on a shared state with local dimension at least d [BYJK08, WCD08, BBT11]. Thus the upper bound $d(\mathcal{S})$ on the entanglement needed cannot be a universal constant, *i.e.*, its value must depend on some measure of complexity of the game.

Our results. We introduce a nonlocal game H for which the answer to question ($\star$) is negative. While $\omega^*(H) = 1$, we show that this success probability can only be achieved in the limit of strategies using entangled states of increasing dimension. More precisely, the game H is such that for any $\varepsilon > 0$,

- H can be won with probability $1 - \varepsilon$ using a shared state of local dimension $O(1/\varepsilon^{7.3\cdots})$;
- any strategy that wins H with probability $1 - \varepsilon$ uses a state of local dimension $\Omega(1/\sqrt{\varepsilon})$.

See Lemma 4 and Theorem 5 for precise statements. The game H is inspired by Hardy's paradox [Har92, Har93]; it has two questions per party and countably infinite answer sets (see Section 3.1 for a brief review of Hardy's paradox, and Section 3.2 for a complete description of the game). We also show that in order to win H with probability $1 - \varepsilon$ the quantum players must use a strategy that assigns positive probability to at least $\Omega(1/\sqrt{\varepsilon})$ distinct answers per party (see Theorem 7). Finally, starting from any two-question game G with finite answer set, we construct games G^δ_∞ with the same features as the Hardy's game, as long as $0 < \delta < \omega(G) - \omega^*(G)$ (see Section 3.4).

Our result has some bearing on the computational complexity of computing the entangled value $\omega^*(G)$ of a general nonlocal game G. If the input size is defined to be the total number of questions and answers in G, then $\omega^*(G)$ is known to be NP-hard to compute [IKM09]. The problem is not known to be in NP however. In fact, no non-trivial upper bounds on its complexity are currently known—it is even unknown whether it can be decided if $\omega^*(G) = 1$. The only known decidability result is for a related parameter $\omega^*_{FV}(G)$, the field-theoretic value of G. Here the tensor-product condition on the players' strategy is relaxed to a commutativity requirement (due to this relaxation, $\omega^*(G) \leq \omega^*_{FV}(G)$ for any game G). The parameter ω^*_{FV} is computable provided that the optimal commuting strategy is finite-dimensional [WCD08, NPA08]. In such a case $\omega^*(G) = \omega^*_{FV}(G)$ since any finite-dimensional commuting strategy can be used to construct a tensor-product strategy with the same success probability. Equality in the general case would follow from a positive resolution of Tsirelson's conjecture (see *e.g.* [Fri12] for a formulation and discussion of the conjecture).

Our results concerning the game H show that no a priori upper bound on the dimension of optimal strategies exists for games with finite question and countable answer sets. Hence, in this case

the "naïve" algorithm that computes $\omega^*(G)$ by performing an exhaustive search over all possible strategies in increasing dimension (and to within increasing precision) will keep finding strategies with increasing success probability. Of course, in practice one will be content with a good approximation to $\omega^*(G)$. In this case one can interpret our result as placing a lower bound, depending on the desired approximation, on the dimension in which the search must be performed. The game H can thus be used as a "dimension witness", *for any dimension*: strategies achieving success probability at least $\omega^*(H) - \varepsilon$ must necessarily use entanglement of local dimension $d = \Omega(1/\sqrt{\varepsilon})$. Examples of similar constructions are already known [BPA+08, Slo11, BBT11], but they all require using different games in order to witness increasing dimensions. In particular, for all the known constructions the number of questions increases along with the dimension to be certified. For instance, Slofstra [Slo11] provides an n-question two-answer XOR game G_n for which achieving $\omega^*(G) - \varepsilon$ requires dimension $\min(2^{\Omega(\sqrt{n})}, \Omega(\varepsilon^{-1/12}))$. Briët et al. [BBT11] provide an $(1/\varepsilon)^{\Theta(1/\varepsilon)}$-question two-answer XOR game G_ε for which the dependence on ε is $\Omega(1/\varepsilon)$. When compared to these constructions, our game has the drawback of having infinitely many possible answers. To rectify this, we can limit our game to answers of length $\ell = \Theta(\log 1/\varepsilon)$. This gives a family of games H_ℓ with two questions and $\mathrm{poly}(1/\varepsilon)$ answers. To win H_ℓ with probability $\omega^*(H_\ell) - \varepsilon$, a state of dimension at least $\Omega(1/\sqrt{\varepsilon})$ is needed (see Corollary 8).

Related work. Prior to our work two examples of nonlocal games were known for which the optimal success probability can only be approached in the limit of infinite-dimensional shared entanglement. Leung et al. [LTW13] consider a game with quantum questions and answers, which has inspired Regev and Vidick [RV13] to produce a game with quantum questions but classical answers. The latter game has two possible (3×3)-dimensional states as questions, and two possible answers per player. Our game provides the first example with classical questions and answers, thus resolving a question first formally asked in [NPA08]. The same question was also asked [WCD08], but specifically for games with finite question and answer sets. Although there is some numerical evidence that infinite-dimensional entanglement may be needed in that case as well [PV10], the question remains tantalizingly open.

Organization of the paper. The remainder of the paper is organized as follows. In Section 2 we formally define the classical and entangled values of a nonlocal game. In Section 3 we review Hardy's paradox, introduce the game H, and show that the classical value of H is $3/4$ while its entangled value equals 1. We also give a more general construction which yields a game with similar properties to Hardy's, but starting from any game with two questions per player and finite answer sets. In Section 4 we establish our main result by showing that restricting the dimension of shared entanglement bounds the achievable success probabilities away from one (see Theorem 5). We also show that restricting the number of different answers has the same effect (see Theorem 7). We conclude in Section 5.

2 Preliminaries

In this section we explain the terminology used to discuss one-round two-player nonlocal games. A reader familiar with nonlocal games is encouraged to proceed directly to the next section.

A two-party nonlocal game G consists of a probability distribution π over a set of the form $\mathcal{I}_A \times \mathcal{I}_B$, called the set of questions; sets of answers $\mathcal{O}_A$ and $\mathcal{O}_B$, and a verification function $V : \mathcal{O}_A \times \mathcal{O}_B \times \mathcal{I}_A \times \mathcal{I}_B \to \{0, 1\}$. With probability $\pi(x, y)$ the referee sends the two players,

traditionally called Alice and Bob, questions $x \in \mathcal{I}_A$ and $y \in \mathcal{I}_B$, respectively. Without communicating, the players must produce answers $a \in \mathcal{O}_A$ and $b \in \mathcal{O}_B$, respectively. They win if $V(a, b, x, y) =: V(a, b|x, y) = 1$ and lose otherwise. Classical players can use shared randomness to possibly enhance their strategy while quantum players can use shared entanglement. The goal of the players is to maximize their probability of winning. In case of classical players, we call this probability the *(classical) value* of game G, denoted $\omega(G)$. In the quantum case we call it the *entangled value* of G, denoted $\omega^*(G)$. Since quantum players are at least as powerful as classical ones, $\omega(G) \leq \omega^*(G)$ for any game G.

2.1 Classical strategies and value

Any classical strategy for a game G can be specified using a pair of functions,

$$\alpha : \mathcal{I}_A \times R \to \mathcal{O}_A \quad \text{and} \quad \beta : \mathcal{I}_B \times R \to \mathcal{O}_B, \tag{1}$$

where R consists of the possible values of shared randomness (which, without loss of generality, also includes any private randomness). We define $\alpha(x, r)$ to be Alice's answer on question x given that the shared randomness takes value r. We define $\beta(y, r)$ similarly. If the shared randomness is distributed according to $\tau : R \to [0, 1]$, then the classical value of the game is

$$\omega(G) := \sup_{\alpha,\beta} \omega(G|\alpha, \beta) := \sup_{\alpha,\beta} \sum_r \tau(r) \sum_{x,y} \pi(x, y) V\Big(\alpha_r(x), \beta_r(y)|x, y\Big), \tag{2}$$

where $\alpha_r(x) := \alpha(x, r)$ and $\beta_r(y) := \beta(y, r)$. Note that for any $r \in R$ the pair (α_r, β_r) specifies a deterministic strategy. Since $\omega(G|\alpha, \beta)$ is a convex combination of $\omega(G|\alpha_r, \beta_r)$, there exists some $r \in R$ for which $\omega(G|\alpha_r, \beta_r) \geq \omega(G|\alpha, \beta)$. Therefore, it suffices to consider only deterministic strategies (α, β), *i.e.*,

$$\omega(G) = \sup_{\substack{\alpha:\mathcal{I}_A\to\mathcal{O}_A \\ \beta:\mathcal{I}_B\to\mathcal{O}_B}} \sum_{x,y} \pi(x, y) V\big(\alpha(x), \beta(y)|x, y\big). \tag{3}$$

The following two lemmas show that if a game either has finitely many questions, or finitely many answers, then the classical value can always be achieved exactly, *i.e.*, the supremum in (3) can be replaced with a maximum.

Lemma 1. *Let $G = (V, \pi)$ be a two-party game with finite question sets $\mathcal{I}_A$, $\mathcal{I}_B$ and (arbitrary) answer sets $\mathcal{O}_A, \mathcal{O}_B$. Then the classical value of G is always achieved by some deterministic strategy, i.e., $\omega(G) = \omega(G|\alpha, \beta)$ for some $\alpha : \mathcal{I}_A \to \mathcal{O}_A$ and $\beta : \mathcal{I}_B \to \mathcal{O}_B$.*

Proof. Consider a two-party game $G = (V, \pi)$ with finite question sets. We have already argued that it suffices to consider only deterministic strategies to find the classical value $\omega(G)$. Therefore, it only remains to show that some deterministic strategy (α, β) achieves success probability $\omega(G)$, *i.e.*, $\omega(G) = \omega(G|\alpha, \beta)$. Since $V(a, b|x, y) \in \{0, 1\}$ and the sets $\mathcal{I}_A, \mathcal{I}_B$ are finite, the sum in Equation (3) can only take on a finite number of different values, each of which is attained by a particular pair of strategies. Therefore, the supremum can be replaced with the maximum and the value of G is achieved by some deterministic strategy. □

Lemma 2. *Let $G = (V, \pi)$ be a 2-party game with countably infinite question sets $\mathcal{I}_A$, $\mathcal{I}_B$ and finite answer sets $\mathcal{O}_A, \mathcal{O}_B$. Then the classical value of G is always achieved by some deterministic strategy.*

Proof. Consider a game $G = (V, \pi)$ with question sets $\mathcal{I}_A, \mathcal{I}_B = \{1, 2, \ldots\}$ and finite answer sets $\mathcal{O}_A, \mathcal{O}_B$. We can identify any deterministic strategy (α, β) with an infinite word

$$\big(\alpha(1), \beta(1)\big)\big(\alpha(2), \beta(2)\big)\big(\alpha(3), \beta(3)\big)\cdots \tag{4}$$

over the finite alphabet $\mathcal{O}_A \times \mathcal{O}_B$. Consider a sequence of deterministic strategies (α_n, β_n) with success probabilities converging to $\omega(G)$. Let w_n be the n-letter prefix of the word corresponding to strategy (α_n, β_n) and consider the set of words $L = \{w_n\}_n$. König's lemma states that if X is an infinite prefix-closed set of words over a fixed alphabet, then we can find an infinite word x whose any prefix is shared by some word in X (see *e.g.* Proposition 1.2.3 in [Lot02]). Since the set L is prefix-closed by construction, there exists w whose any prefix is shared by some word in L. It now remains to note that the strategy that corresponds to w achieves success probability $\omega(G)$. This is because the partial strategies corresponding to prefixes of w achieve success probabilities arbitrarily close to $\omega(G)$. □

In the case when both the question sets and answer sets are allowed to be infinite the value of the game cannot always be achieved. To see this consider the following game, which was suggested to us by A. Belovs. The question and answer sets are $\mathcal{I}_A, \mathcal{I}_B, \mathcal{O}_A, \mathcal{O}_B := \mathbb{N}$, $\pi(x, y) = 2^{-x-y}$ and $V(a, b|x, y) = 1$ if and only if $a, b > \max(x, y)$. Clearly no deterministic strategy achieves success probability one. Yet if we define (α_n, β_n) via $\alpha(x) = x + n$ and $\beta_n(y) = y + n$, then (α_n, β_n) gives a sequence of strategies with vanishing error. Hence, $\omega(G) = 1$ even though no single strategy can be used to win game G perfectly.

2.2 Quantum strategies and value

A quantum strategy is specified by a finite-dimensional shared state $|\psi\rangle \in \mathbb{C}^{d_A} \otimes \mathbb{C}^{d_B}$, a POVM $\mathcal{A}^x = \{A_a^x : a \in \mathcal{O}_A\}$ for every question $x \in \mathcal{I}_A$ to Alice, and a POVM $\mathcal{B}^y = \{B_b^y : b \in \mathcal{O}_B\}$ for every question $y \in \mathcal{I}_B$ to Bob. Upon receiving questions x and y, the parties measure their systems with POVMs $\mathcal{A}^x$, $\mathcal{B}^y$ respectively and answer with their respective measurement outcomes a and b. The dimension of a strategy is the dimension of the shared entangled state it employs. We also use the term "local dimension" to refer to d_A or d_B, *i.e.*, the dimension of Alice's or Bob's local systems.

The quantum value of the game is given by

$$\omega^*(G) = \sup \sum_{x,y} \pi(x, y) V(a, b|x, y) \langle\psi| \big(A_a^x \otimes B_b^y\big) |\psi\rangle, \tag{5}$$

where the supremum is taken over all finite-dimensional shared states $|\psi\rangle$ and sets of POVMs $\{\mathcal{A}^x\}_x$ and $\{\mathcal{B}^y\}_y$.

3 Hardy's game and its generalization

We introduce a nonlocal game that we call Hardy's game, as it is based on Hardy's paradox [Har92, Har93]. We first describe the paradox, then the game we derive from it, and exhibit a sequence of good strategies for quantum players of this game. We end this section by giving a general construction which yields a nonlocal game with similar properties as Hardy's game, starting from any game G with finite question and answer sets.

3.1 Hardy's paradox

Hardy's paradox is a two-party experimental setup whose outcomes, as predicted by quantum mechanics, cannot be reproduced by any local hidden variables theory. The setup is noteworthy for being minimal in a sense that it involves only two qubits, on which both of the parties perform two measurements with two outcomes each. Our description of the setup roughly follows the explanations given in [Lal01].

Suppose Alice and Bob share a two-qubit state

$$|\psi\rangle := \frac{\sin(\theta)|v'_0, v'_0\rangle - \cos(\theta)\big(|v'_0, v'_1\rangle + |v'_1, v'_0\rangle\big)}{\sqrt{1+\cos^2(\theta)}}, \tag{6}$$

where $\mathcal{S}' = \{|v'_0\rangle, |v'_1\rangle\} \subseteq \mathbb{R}^2$ is any orthonormal basis and $0 < \theta < \frac{\pi}{2}$. Let $\mathcal{S} = \{|v_0\rangle, |v_1\rangle\} \subseteq \mathbb{R}^2$ be the basis obtained by rotating $\mathcal{S}'$ by angle 2θ in the xz plane of the Bloch sphere, *i.e.*,

$$|v_0\rangle = \cos(\theta)|v'_0\rangle + \sin(\theta)|v'_1\rangle, \tag{7}$$
$$|v_1\rangle = -\sin(\theta)|v'_0\rangle + \cos(\theta)|v'_1\rangle. \tag{8}$$

If Alice and Bob measure the state $|\psi\rangle$ using the bases $\mathcal{S}$ or $\mathcal{S}'$ the following conditions are satisfied (see also Figure 1):

(i) if both parties measure in $\mathcal{S}$, the outcome pair $(0,0)$ occurs with probability $p_\theta > 0$ (see Equation (9) for an exact formula);

(ii) if one of the parties measures in $\mathcal{S}$ and the other in $\mathcal{S}'$, the outcome pair $(0,0)$ never occurs;

(iii) if both parties measure in $\mathcal{S}'$, the outcome pair $(1,1)$ never occurs.

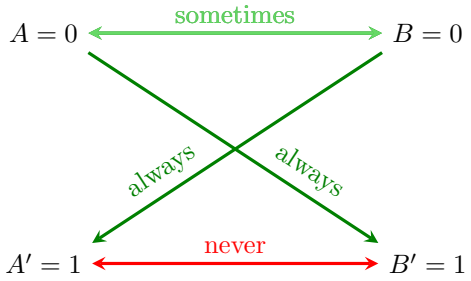

Figure 1: Summary of correlations between the outcomes of Alice and Bob's measurements in Hardy's setup. For instance, the arrow $A = 0 \to B' = 1$ marked "always" means that, if Alice measures her system in basis $\mathcal{S}$ and Bob measures his system in basis $\mathcal{S}'$, then whenever Alice obtains the outcome 0 Bob always obtains the outcome 1.

To see that condition (i) is satisfied, we note that

$$p_\theta := |\langle v_0, v_0|\psi\rangle|^2 = \frac{\cos^4(\theta)\sin^2(\theta)}{1+\cos^2(\theta)} > 0 \tag{9}$$

as $0 < \theta < \pi$. To verify the other two conditions one can check that $\langle v_0, v'_0|\psi\rangle = 0$, $\langle v'_0, v_0|\psi\rangle = 0$, and $\langle v'_1, v'_1|\psi\rangle = 0$.

Any two-qubit state that is neither product nor maximally entangled can be expressed in the form of Equation (6) for appropriately chosen basis $\mathcal{S}'$ and angle θ [Har93, Gol94]. Therefore, almost any two-qubit state can be used to perform an experiment whose outcomes will satisfy

conditions (i)–(iii). Different values of θ will result in different p_θ; one can verify that the maximum p_θ is $\frac{5\sqrt{5}-11}{2} \approx 0.09$ and is attained at $\theta = \arccos\left(\left(\frac{\sqrt{5}-1}{2}\right)^{1/2}\right) < \frac{\pi}{4}$.

Local hidden variables theories cannot reproduce the predictions (i)–(iii) made by quantum mechanics. Any such theory assigns definite outcomes (depending only on the hidden variables) for each of the four measurements: A, A', B and B'. Condition (i) guarantees that for some setting of the hidden variables the outcomes associated with A and B will both be 0. According to Condition (ii) this implies that for the same setting of hidden variables the outcomes associated with B' and A' must both be 1. This, however, contradicts Condition (iii).

A different construction that is often used to disprove the existence of non-contextual hidden variables theories are the so-called Kochen-Specker sets [Gle57, KS67]. These sets can easily be turned into a nonlocal game with entangled value 1 and classical value bounded away from 1 [CHTW04]. Hardy's paradox in itself does not immediately yield such a game; indeed it is unclear how to enforce Condition (i) in the framework of nonlocal games. In the next section we show how to accommodate with that condition and turn Hardy's paradox into a nonlocal game.

3.2 A nonlocal game from Hardy's paradox

Building on Hardy's paradox, we now propose the following nonlocal game that we call *Hardy's game* and denote as H. In game H there are two questions per player: Alice's question set is $\mathcal{I}_A = \{A, A'\}$, Bob's question set is $\mathcal{I}_B = \{B, B'\}$, and the questions are uniformly distributed, *i.e.*, $\pi(X, Y) = \frac{1}{4}$ for all $X \in \mathcal{I}_A$, $Y \in \mathcal{I}_B$. The possible answers for both parties are binary strings of arbitrary but finite length, *i.e.*, $\mathcal{O}_A, \mathcal{O}_B = \{0,1\}^*$. The verification function $V : \mathcal{O}_A \times \mathcal{O}_B \times \mathcal{I}_A \times \mathcal{I}_B \to \{0,1\}$ is defined by $V(a, b|X, Y) = 1$ if and only if all of the following conditions are satisfied.

1. The answer strings a and b have the same length, *i.e.*, $|a| = |b|$.
2. If $X = A$ and $Y = B$, then $a_i = b_i = 0$ for some position $i \in [n]$, where $n := |a| = |b|$.
3. For each position $i \in [n]$:
   (a) if $(X, Y) = (A, B')$ or $(X, Y) = (A', B)$, then $a_i = 1$ or $b_i = 1$;
   (b) if $(X, Y) = (A', B')$, then $a_i = 0$ or $b_i = 0$.

In the above, Condition 3 requires that each pair of answer bits (a_i, b_i) satisfies the last two conditions in Hardy's paradox. Condition 2 requires that there exists a position i for which (a_i, b_i) satisfy the first condition in Hardy's paradox with certainty.

3.3 The classical and entangled values

In this section we determine both the classical value and the entangled value of Hardy's game.

Lemma 3. *The classical value of Hardy's game is* $\omega(H) = 3/4$.

Proof. The value of Hardy's game is given by

$$\omega(H) = \max_{\alpha,\beta} \sum_{X \in \{A,A'\}} \sum_{Y \in \{B,B'\}} \frac{1}{4} V\big(\alpha(X), \beta(Y)|X, Y\big) \tag{10}$$

and thus $\omega(H) \in \{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\}$. In fact, $\omega(H) \geq \frac{3}{4}$, since a strategy defined via $\alpha(A) = \beta(B) = 0$ and $\alpha(A') = \beta(B') = 1$ achieves a success probability of $\frac{3}{4}$. To complete the proof it remains to rule out the existence of a perfect deterministic strategy. For contradiction, assume that $\omega(H|\alpha, \beta) = 1$ for some deterministic strategy (α, β). Since the strategy is perfect, there exists i such that the

ith bit of the strings $\alpha(A)$ and $\beta(B)$ is zero. This implies that the ith bit of both strings $\alpha(A')$ and $\beta(B')$ equals one. Therefore, the players lose on question (A', B') which gives the desired contradiction. □

Lemma 4. *The entangled value of Hardy's game is $\omega^*(H) = 1$. Furthermore, for any $\varepsilon > 0$ there exists a quantum strategy that succeeds with probability at least $1-\varepsilon$ using only answers of length $\ell = \Theta(\log(1/\varepsilon))$ and entangled state of local dimension $d = \Theta(\varepsilon^{1/\log c}) = \Theta(1/\varepsilon^{7.35\cdots})$, where $c = \frac{13-5\sqrt{5}}{2}$.*

Proof. For each $n \in \mathbb{N}$ consider a strategy $\mathfrak{S}_n$ in which players share n copies of the state $|\psi\rangle$ defined in Equation (6), for some value of θ. To produce an n-bit answer string they measure each of the n copies in basis $\mathcal{S} = \{|v_0\rangle, |v_1\rangle\}$ upon receiving an unprimed question and in basis $\mathcal{S}'$ (see Equation (8)) upon receiving a primed question. To analyze the success probability of strategy $\mathfrak{S}_n$ recall the three conditions from Section 3.2. Since both players answer strings of length n, Condition 1 is always satisfied. As discussed in Section 3.1, the outcomes obtained by measuring $|\psi\rangle$ in basis $\mathcal{S}$ and $\mathcal{S}'$ satisfy the conditions from Hardy's paradox. This implies that Alice's and Bob's answer bits (a_i, b_i) always satisfy Condition 3. Finally, note that upon question (A, B) we have $a_i = b_i = 0$ with probability p_θ (see Equation (9)). Hence, Condition 2 is satisfied with probability $1 - (1 - p_\theta)^n$. If $\theta = \arccos\left(\left(\frac{\sqrt{5}-1}{2}\right)^{1/2}\right)$, then $\mathfrak{S}_n$ errs with probability $\frac{1}{4}\left(\frac{13-5\sqrt{5}}{2}\right)^n$. Let $c = \frac{13-5\sqrt{5}}{2} \approx 0.91 < 1$, fix any $\varepsilon > 0$ and consider $n = \left\lceil \frac{\log \varepsilon}{\log c} \right\rceil = \Theta(\log(1/\varepsilon))$. Then $\mathfrak{S}_n$ errs with probability at most ε, answers strings of length $\Theta(\log(1/\varepsilon))$ and uses entanglement of local dimension $d = 2^n = \Theta(\varepsilon^{1/\log c})$. □

3.4 A general construction

The construction we used to define Hardy's game from Hardy's paradox can also be applied when starting from any two-player game G with finite question and answer sets. For any such game $G = (V, \pi)$ and any $\delta > 0$ we define a game $G^\delta_\infty = (V^\delta_\infty, \pi)$. Results similar to those in Section 3.3 can be obtained for any of the defined games. Yet we focus on the games G^δ_∞ arising from games G with two questions per player, denoted A, A' and B, B' respectively, and $\delta < \omega^*(G) - \omega(G)$, since our main contribution in Section 4 only applies in this case. For example, our construction can be applied to the CHSH game, to obtain $\mathrm{CHSH}^\delta_\infty$ for any $0 < \delta < \cos^2(\frac{\pi}{8}) - \frac{3}{4}$.

Fix some entangled strategy that succeeds at G with probability $\omega^*(G) - \frac{\delta}{2}$ and let ω^*_{XY} be its success probability conditioned on questions (X, Y) being asked. The construction of the game G^δ_∞ is very simple. In the game, the verifier first chooses questions $(X, Y) \in \{A, A'\} \times \{B, B'\}$ according to π, and sends X to Alice and Y to Bob. Each player may reply with a string of arbitrary yet finite length over the respective answer alphabets in G, *i.e.*, $\mathcal{O}_A$ and $\mathcal{O}_B$. The players win, *i.e.*, $V^\delta_\infty(a, b|X, Y) = 1$, if and only if their respective answer strings $a = (a_i) \in \mathcal{O}^*_A$ and $b = (b_i) \in \mathcal{O}^*_B$ have the same length and the fraction of positions i for which $V(a_i, b_i|X, Y) = 1$ is at least $\omega^*_{XY} - \frac{\delta}{2}$. The above definition of G^δ_∞ depends on the probabilities ω^*_{XY} and hence one should rather write $G^\delta_\infty(\omega^*_{00}, \omega^*_{10}, \omega^*_{01}, \omega^*_{11})$. For the sake of simplicity from now on we concentrate on the case where it is possible to choose a strategy for which the success probability does not depend on the question (X, Y), *i.e.*, $\omega^*_{XY} = \omega^*(G) - \frac{\delta}{2}$.

Our results for Hardy's game in Section 3.3 have straightforward extensions to the general case. More precisely, for any game G with finite question and answer sets and any $\delta > 0$, we have $\omega(G^\delta_\infty) = \omega(G)$, $\omega^*(G^\delta_\infty) = 1$ and the players can achieve success probability $1 - \varepsilon$ using entanglement of local dimension $\mathrm{poly}\left(\frac{1}{\varepsilon}\right)$.

4 Finite strategies do not achieve the entangled value

The results in this section apply to Hardy's game as well as any game G_∞^δ which is constructed from some game G which has *two questions* per player and finite answer sets $\mathcal{O}_A$ and $\mathcal{O}_B$ as explained in Section 3.4. (We suspect that this restriction can be relaxed but did not investigate this further.) We show that for all small enough $\delta > 0$ the entangled value cannot be attained by any finite-dimensional strategy. We also show that the same holds for strategies with bounded answer length (irrespective of the dimension of the entangled state they are based on).

Intuitively, the need for unbounded amount of entanglement to successfully play G_∞^δ stems from the requirement that the answer strings in G_∞^δ should on average achieve the value of at least $\omega^*(G) - \frac{\delta}{2}$. In order to guarantee success the players are "required" to repeatedly sample from an entangled strategy for G, necessitating more and more entanglement as they improve their chances by exploiting concentration.

At a slightly more formal level, our proof uses arguments that are reminiscent of techniques used to lower bound the amount of quantum communication required to send classical information from Alice to Bob in quantum one-way communication setting. We show that increasingly better quantum strategies for G_∞^δ give rise to residual states (post-measurement states on Bob's side after Alice's measurement) that can be distinguished with increasing success probability. We obtain the desired bound on the dimension by observing that the latter can only happen if the states are of sufficiently high dimension. The precise argument is slightly more involved as indeed success in the game does not imply any form of communication, and we cannot use the intuition directly. Details are given in the proof of Theorem 5 in the next section.

4.1 Strategies with bounded entanglement

Theorem 5. *Let $\varepsilon > 0$. Consider a strategy which succeeds at Hardy's game H with probability at least $1 - \varepsilon$ and uses an entangled state of local dimensions d_A and d_B respectively. Then $\min(d_A, d_B) \geq 1/(16\sqrt{\varepsilon})$. As a consequence, the game H cannot be won with probability one using finite-dimensional entanglement.*

Furthermore, a similar statement applies for any game G_∞^δ with $0 < \delta < \omega^*(G) - \omega(G)$, constructed from a two-player game G with two questions per player and finite answer sets. In this case, $\min(d_A, d_B) = O\big(\frac{1}{\sqrt{\varepsilon}}\big)$, where the implied constant depends on $\min_{XY} \pi(X, Y)$ only.[1]

Below we give a proof of the "furthermore" part of Theorem 5 which applies for the games G_∞^δ. Although, strictly speaking, Hardy's game does not fall in this category, the proof is in all points similar. Before proceeding we give a key lemma with the following informal interpretation. Consider any pair of answers (s, t) associated with Alice's questions A and A' respectively in G_∞^δ. Then we can find a question $Y \in \{B, B'\}$ for Bob such that for any $u \in \mathcal{O}_B^*$ either the answer pair (s, u) is rejected upon question (A, Y) or the answer pair (t, u) is rejected upon question (A', Y). To state the lemma, let us define the following sets. For any $s, t \in \mathcal{O}_A^*$ let

$$\mathcal{U}_s := \{u \in \mathcal{O}_B^* : \ V_\infty^\delta(s, u|A, B) = 0, \textit{ i.e.}, (s, u) \text{ is rejected upon questions } (A, B)\} \tag{11}$$

[1] For games with two questions per player $\min_{XY} \pi(X, Y)$ is necessarily positive, unless $\omega^*(G) = \omega(G)$. To see this, suppose G is such that *e.g.* $\pi(A', B') = 0$. Fix an optimal quantum strategy. Consider a classical strategy in which each player, upon receiving the unprimed question, uses shared randomness to sample from the distribution produced by the quantum strategy upon questions (A, B). Now, upon receiving the primed question a player knows that the other player received the unprimed question. Hence the player can use the same shared randomness to again simulate the behavior of the quantum players: for any value of the shared random string the player knows precisely which answer the other player will provide. Hence the classical strategy reproduces the quantum strategy perfectly.

and let $\mathcal{P}_t$ be the set of answers u for Bob such that (t,u) is rejected upon question (A',B). Similarly define $\mathcal{U}'_s$ and $\mathcal{P}'_t$ as the sets of Bob's answers corresponding to rejected answer pairs for questions (A,B') and (A',B') respectively.[2]

Lemma 6. *For any $(s,t) \in (\mathcal{O}_A^* \times \mathcal{O}_A^*)$ it holds that either $\mathcal{U}'_s \cup \mathcal{P}'_t = \mathcal{O}_B^*$ or $\mathcal{U}_s \cup \mathcal{P}_t = \mathcal{O}_B^*$.*

Proof. For contradiction suppose that there exists a pair of answer strings $(s,t) \in (\mathcal{O}_A^* \times \mathcal{O}_A^*)$ such that both $\mathcal{U}'_s \cup \mathcal{P}'_t \subsetneq \mathcal{O}_B^*$ and $\mathcal{U}_s \cup \mathcal{P}_t \subsetneq \mathcal{O}_B^*$. Hence there exists a pair of strings of answers $(u,v) \in (\mathcal{O}_B^* \times \mathcal{O}_B^*)$ for Bob such that $u \notin \mathcal{U}_s \cup \mathcal{P}_t$ and $v \notin \mathcal{U}'_s \cup \mathcal{P}'_t$. By definition, this means that the pair of answers (s,u), (s,v), (t,u) and (t,v) are accepted by the verifier in the game G_∞^δ, on questions (A,B), (A,B'), (A',B) and (A',B') respectively. We now derive a contradiction by defining a classical strategy for the game G that has success probability strictly larger than $\omega(G)$. The strategy is simple: upon receiving their respective questions, the players use shared randomness to select an integer $i \in \{1,\ldots,|s|\}$ (note that necessarily $|s| = |t| = |u| = |v|$) and answer the i-th symbol in their strings s or t for Alice (depending on whether her question was A or A'), and u or v for Bob (depending on whether his question was B or B'). By definition of V_∞^δ, their success probability is at least $\omega^*(G) - \delta > \omega(G)$, a contradiction. □

Based on Lemma 6 one can already give a short proof of the fact that no finite strategy can achieve success probability $\omega^*(G_\infty^\delta) = 1$ for any of the considered games G_∞^δ. To see this, assume that a perfect finite strategy exists and let ρ_s be Bob's residual state in the case when Alice has responded with string s upon question A and let ρ'_t be his state in the case when Alice has responded t upon question A'. Note that if $\mathcal{U}_s \cup \mathcal{P}_t = \mathcal{O}_B^*$ then any answer of Bob's which is accepted upon question (A,B) and Alice answering s would be rejected upon question (A',B) and Alice answering t. Hence, in order to win Bob must be able to perfectly distinguish state ρ_s from ρ_t, *i.e.*, $\rho_s \perp \rho_t$. Lemma 6 allows to conclude that $\rho_s \perp \rho_t$ for all s,t which is a contradiction since $\sum_s \rho_s = \sum_t \rho_t$ cannot be the zero matrix.

To prove the dimension estimate stated in Theorem 5 we need a more involved quantitative argument which we are about to give. Nevertheless, the underlying idea remains the same.

Proof of Theorem 5. Fix $\varepsilon > 0$, and assume that there exists a quantum strategy for winning G_∞^δ with probability $1-\varepsilon$ using $|\psi\rangle \in \mathbb{C}^{d_A} \otimes \mathbb{C}^{d_B}$ as the shared entanglement. Without loss of generality we can assume that $|\psi\rangle$ has full Schmidt rank and $d_A = d_B =: d$. Next, let $\{A_s : s \in \mathcal{O}_A^*\}$ be the POVM measured by Alice upon question A. Define POVM $\{A'_t : t \in \mathcal{O}_A^*\}$ similarly, as well as POVMs $\{B_u : u \in \mathcal{O}_B^*\}$ for Bob on question B and $\{B'_v : v \in \mathcal{O}_B^*\}$ on question B'. Also, for any $s,t \in \mathcal{O}_A^*$, define the following post-measurement states on Bob's system

$$\rho_s := \mathrm{Tr}_A\left((A_s \otimes \mathrm{Id})|\psi\rangle\langle\psi|\right) \quad \text{and} \quad \rho'_t := \mathrm{Tr}_A\left((A'_t \otimes \mathrm{Id})|\psi\rangle\langle\psi|\right), \tag{12}$$

and let $\rho := \sum_s \rho_s = \sum_t \rho'_t = \mathrm{Tr}_A(|\psi\rangle\langle\psi|)$. Recall the definition of the set $\mathcal{U}'_s$ in Equation (11), and define $U'_s := \sum_{v \in \mathcal{U}'_s} B_v$. Similarly define $U_s := \sum_{u \in \mathcal{U}_s} B_u$, $P'_t := \sum_{v \in \mathcal{P}'_t} B_v$ and $P_t := \sum_{u \in \mathcal{P}_t} B_u$. Our assumption that the strategy succeeds with probability $1-\varepsilon$ implies that the following four

[2] The letters $\mathcal{P}$ and $\mathcal{U}$ are meant to stand for the primed and unprimed questions for Alice.

relations must hold:

$$\pi(A, B) \sum_s \mathrm{Tr}(\rho_s U_s) \leq \varepsilon \tag{13}$$

$$\pi(A, B') \sum_s \mathrm{Tr}(\rho_s U'_s) \leq \varepsilon \tag{14}$$

$$\pi(A', B) \sum_t \mathrm{Tr}(\rho'_t P_t) \leq \varepsilon \tag{15}$$

$$\pi(A', B') \sum_t \mathrm{Tr}(\rho'_t P'_t) \leq \varepsilon. \tag{16}$$

To proceed we define the following sets:

$$S' := \{(s,t) \in (\mathcal{O}_A^* \times \mathcal{O}_A^*) : \mathcal{U}'_s \cup \mathcal{P}'_t = \mathcal{O}_B^*\}$$
$$S := \{(s,t) \in (\mathcal{O}_A^* \times \mathcal{O}_A^*) \setminus S' : \mathcal{U}_s \cup \mathcal{P}_t = \mathcal{O}_B^*\}.$$

By the definition of the two sets and Lemma 6, the set $\mathcal{O}_A^* \times \mathcal{O}_A^*$ is the disjoint union of S and S'. We derive a lower bound on d by studying the following quantity:

$$\mathrm{Tr}(\rho^2) = \sum_{s,t \in \mathcal{O}_A^*} \mathrm{Tr}(\rho_s \rho'_t) = \sum_{(s,t) \in S'} \mathrm{Tr}(\rho_s \rho'_t) + \sum_{(s,t) \in S} \mathrm{Tr}(\rho_s \rho'_t). \tag{17}$$

Since ρ is a d-dimensional density matrix it holds that $\mathrm{Tr}(\rho^2) \geq 1/d$. We now upper bound each of the two terms on the right-hand side of Equation (17) as a function of ε. The two bounds are similar; we analyze the summation corresponding to pairs $(s,t) \in S'$.

By the definition of the set S', we have that $\mathcal{U}'_s \cup \mathcal{P}'_t = \mathcal{O}_B^*$ for any $(s,t) \in S'$. Hence we can decompose

$$\mathrm{Id} = U'_s + P'_t - R'_{s,t}, \tag{18}$$

where the term $R'_{s,t} := \sum_{v \in \mathcal{U}'_s \cap \mathcal{P}'_t} B'_v$ takes care of double counting. Write

$$\begin{aligned}\sum_{(s,t) \in S'} \mathrm{Tr}(\rho_s \rho'_t) &= \sum_{(s,t) \in S'} \mathrm{Tr}(\rho_s (U'_s + P'_t - R'_{s,t}) \rho'_t) \\ &= \sum_{(s,t) \in S'} \mathrm{Tr}(\rho_s U'_s \rho'_t) + \sum_{(s,t) \in S'} \mathrm{Tr}(\rho_s P'_t \rho'_t) - \sum_{(s,t) \in S'} \mathrm{Tr}(\rho_s R'_{s,t} \rho'_t).\end{aligned} \tag{19}$$

We bound each of the three terms from Equation (19) separately. The first two can be bounded in a similar way. Specifically, for the first term, if we define $\sigma_s := \sum_{t:\,(s,t) \in S'} \rho'_t$ then

$$\begin{aligned}\sum_{(s,t) \in S'} \mathrm{Tr}(\rho_s U'_s \rho'_t) &\leq \sum_s \left|\mathrm{Tr}(\rho_s U'_s \sigma_s)\right| \\ &\leq \sum_s \Big(\mathrm{Tr}(\rho_s U'_s \sigma_s^2 U'_s)\Big)^{1/2} \Big(\mathrm{Tr}(\rho_s)\Big)^{1/2} \\ &\leq \Big(\sum_s \mathrm{Tr}(\rho_s U'_s \sigma_s^2 U'_s)\Big)^{1/2} \Big(\sum_s \mathrm{Tr}(\rho_s)\Big)^{1/2} \\ &\leq \Big(\sum_s \mathrm{Tr}(\rho_s U'_s)\Big)^{1/2} \\ &\leq \sqrt{\varepsilon / \pi(A, B')},\end{aligned}$$

where the second inequality uses the Cauchy-Schwarz inequality for the trace inner product (using cyclicity of the trace to write $\mathrm{Tr}(\rho_s U'_s \sigma_s) = \mathrm{Tr}(\sqrt{\rho_s} U'_s \sigma_s \cdot \sqrt{\rho_s})$), the fourth uses $\sigma_s^2 \leq \mathrm{Id}$, $(U'_s)^2 \leq U'_s$, and $\mathrm{Tr}(\rho) \leq 1$, and the last is by (14). The same bound can be derived for the second term in Equation (19), this time using (16). Note that the last sum in Equation (19) may be negative, and we need to bound it as well. We proceed as follows:

$$\begin{aligned}
\Big| \sum_{(s,t)\in S_1} \mathrm{Tr}(\rho_s R'_{s,t} \rho'_t) \Big| &\leq \Big(\sum_{(s,t)\in S'} \mathrm{Tr}(\rho_s R'_{s,t} \rho'_t R'_{s,t}) \Big)^{1/2} \Big(\sum_{(s,t)\in S_1} \mathrm{Tr}(\rho_s \rho'_t) \Big)^{1/2} \\
&\leq \Big(\sum_{(s,t)\in S_1} \mathrm{Tr}\left(\sqrt{R'_{s,t}} \rho_s \sqrt{R'_{s,t}}\right) \mathrm{Tr}\left(\sqrt{R'_{s,t}} \rho'_t \sqrt{R'_{s,t}}\right) \Big)^{1/2} \Big(\mathrm{Tr}(\rho^2) \Big)^{1/2} \\
&\leq \Big(\sum_{(s,t)\in S_1} \mathrm{Tr}(\rho_s U'_s) \, \mathrm{Tr}(\rho'_t P'_t) \Big)^{1/2} \\
&\leq \Big(\Big(\sum_s \mathrm{Tr}(\rho_s U'_s) \Big)^{1/2} \Big(\sum_t \mathrm{Tr}(\rho'_t P'_t) \Big) \Big)^{1/2} \\
&\leq \varepsilon / \sqrt{\pi(A,B')\pi(A',B')},
\end{aligned}$$

where the first inequality uses the Cauchy-Schwarz inequality for the trace inner product (using cyclicity of the trace to write $\mathrm{Tr}(\rho_s R'_{s,t} \rho'_t) = \mathrm{Tr}(\rho_s^{1/2} R'_{s,t} (\rho'_t)^{1/2} \cdot (\rho'_t)^{1/2} \rho_s^{1/2})$), the second uses the inequality $\mathrm{Tr}(XY) \leq \mathrm{Tr}(X)\,\mathrm{Tr}(Y)$ valid for all positive X, Y, and the fact that $\mathrm{Tr}(\rho_s \rho'_t) \geq 0$ for every (s,t), the third uses cyclicity of the trace and $R'_{s,t} \leq U'_s$, $R'_{s,t} \leq P'_t$, the fourth uses that all terms are non-negative to upper bound the summation over $(s,t) \in S'$ by a summation over all s, t, and the last uses (14) and (16).

Putting everything together, we have shown that $\big| \sum_{(s,t)\in S'} \mathrm{Tr}(\rho_s \rho_t) \big| = O(\sqrt{\varepsilon})$, where the implied constant depends only on $\min_{XY} \pi(X,Y)$. Proceeding in a similar way, the same bound can be obtained for the summations over pairs $(s,t) \in S$. From Equation (17) we then obtain

$$\frac{1}{d} \leq \sum_{s,t} \mathrm{Tr}(\rho_s \rho'_t) = O\big(\sqrt{\varepsilon}\big),$$

proving the bound claimed in the theorem. □

4.2 Strategies with bounded number of answers

Theorem 7. *Let G be a two-player game with two questions per player and finite answer sets $\mathcal{O}_A$ and $\mathcal{O}_B$. Also let $\varepsilon > 0$ and $0 < \delta < \omega^*(G) - \omega(G)$. Consider an entangled strategy that succeeds at G_∞^δ with probability at least $1 - \varepsilon$ and for which either Alice's or Bob's answers always fall within a set of cardinality L. Then $L = \Omega(1/\sqrt{\varepsilon})$, where the implied constant depends on $\min_{XY} \pi(X,Y)$ only.*

Proof. Let $L \in \mathbb{N}$ and fix a string $s \in \mathcal{O}_A^*$ of length at most $\log(L)$ that is answered by Alice on question A. Let σ_0 be the post-measurement state on Alice's system that corresponds to Bob applying his POVM $\{B_u\}$ upon question B and obtaining an answer string $u \notin \mathcal{U}_s$, that is,

$$\sigma_0 := \sum_{u \notin \mathcal{U}_s} \mathrm{Tr}_A \left((\mathrm{Id} \otimes B_u) |\psi\rangle\langle\psi| \right). \tag{20}$$

If the questions are (A, B) and Alice obtains s as outcome, then by definition of the set $\mathcal{U}_s$ all other answers for Bob are rejected in the game, hence

$$\mathrm{Tr}\left(A_s(\rho - \sigma_0)\right) = \mathrm{Tr}(\rho_s) - \mathrm{Tr}(A_s \sigma_0) \leq \varepsilon / \pi(A,B), \tag{21}$$

where $\{A_s : s \in \mathcal{O}_A^*\}$ is the POVM measured by Alice upon the unprimed question A. Similarly, let σ_1' be the post-measurement state on Alice's system that corresponds to Bob applying his POVM for question B' and obtaining an answer string $v \notin \mathcal{U}_s'$. By the same reasoning,

$$\operatorname{Tr}\left(A_s(\rho - \sigma_1')\right) \;=\; \operatorname{Tr}(\rho_s) - \operatorname{Tr}(A_s\sigma_1') \;\leq\; \varepsilon/\pi(A, B'). \tag{22}$$

Let P' be the sum of those POVM elements for Alice's question A' that correspond to answers t such that $\mathcal{U}_s' \cup \mathcal{P}_t' = \mathcal{O}_B^*$, and $Q' = \mathrm{Id} - P'$. By Lemma 6, every POVM element in Q' corresponds to an answer t such that $\mathcal{U}_s \cup \mathcal{P}_t = \mathcal{O}_B^*$. On questions (A', B), if Alice obtains one of the outcomes t corresponding to P' and Bob obtains an outcome $u \notin \mathcal{U}_s'$ then it must be that $u \in \mathcal{P}_t$ and hence $\operatorname{Tr}(P'\sigma_1') \leq \varepsilon/\pi(A', B)$. Similarly, on questions (A', B') if if Alice obtains one of the outcomes t corresponding to Q' and Bob obtains an outcome $u \notin \mathcal{U}_s$ then $u \in \mathcal{P}_t$ and hence $\operatorname{Tr}(Q'\sigma_0) \leq \varepsilon/\pi(A', B')$. We have proven:

$$\operatorname{Tr}(Q'\sigma_0) \leq \varepsilon/\pi(A', B') \qquad \text{and} \qquad \operatorname{Tr}(P'\sigma_1') \leq \varepsilon/\pi(A', B). \tag{23}$$

Using $P' + Q' = \mathrm{Id}$ we may decompose

$$\begin{aligned}
\operatorname{Tr}(A_s\sigma_0) &= \operatorname{Tr}(A_sP'\sigma_0) + \operatorname{Tr}(A_sQ'\sigma_0) \\
&\leq \operatorname{Tr}(A_sP'\sigma_0) + \operatorname{Tr}(A_sQ'A_s\sigma_0)^{1/2} \operatorname{Tr}(Q'\sigma_0)^{1/2} \\
&\leq \operatorname{Tr}(A_sP'\sigma_0) + \sqrt{\varepsilon/\pi(A', B')} \\
&= \operatorname{Tr}(A_sP'\rho) - \operatorname{Tr}(A_sP'(\rho - \sigma_0)) + \sqrt{\varepsilon/\pi(A', B')} \\
&\leq \operatorname{Tr}(A_sP'\rho) + \operatorname{Tr}(P'A_sP'(\rho - \sigma_0))^{1/2} \operatorname{Tr}(A_s(\rho - \sigma_0))^{1/2} + \sqrt{\varepsilon/\pi(A', B')} \\
&\leq \operatorname{Tr}(A_sP'\rho) + O\big(\sqrt{\varepsilon}\big),
\end{aligned} \tag{24}$$

where the first inequality is Cauchy-Schwarz, the second uses $A_sQ'A_s \leq \mathrm{Id}$ and (23), the third inequality is Cauchy-Schwarz again, and the last uses $P'A_sP' \leq \mathrm{Id}$ and (21). A similar chain of inequalities leads to the bound

$$\operatorname{Tr}(A_s\sigma_1') \leq \operatorname{Tr}(A_sQ'\rho) + O\big(\sqrt{\varepsilon}\big). \tag{25}$$

But

$$\operatorname{Tr}(A_sP'\rho) + \operatorname{Tr}(A_sQ'\rho) = \operatorname{Tr}(A_s\rho) = \operatorname{Tr}(\rho_s),$$

which combined with (24), (25) and (21), (22) yields

$$2\operatorname{Tr}(\rho_s) \;\leq\; \operatorname{Tr}(A_s\sigma_0) + \operatorname{Tr}(A_s\sigma_1') + O\big(\sqrt{\varepsilon}\big) \;\leq\; \operatorname{Tr}(\rho_s) + O\big(\sqrt{\varepsilon}\big),$$

implying that $\operatorname{Tr}(\rho_s) = O\big(\sqrt{\varepsilon}\big)$. Summing over s we deduce $1 \leq O(\sqrt{\varepsilon})L$, proving the theorem. □

We can use Hardy's game, or any of the games G_∞^δ, to certify the dimension of an entangled state. According to Theorem 5, if Alice and Bob succeed with high probability, then they must possess an entangled state of high local dimension. In practice it is not reasonable to test entanglement via a scheme involving infinitely many outcomes. Therefore, we consider games $G_\infty^\delta(\ell)$ that are obtained from G_∞^δ by limiting the parties to answer strings of length at most $\ell \in \mathbb{N}$. For an appropriately chosen $\ell = \Theta(\log 1/\varepsilon)$, the game $G_\infty^\delta(\ell)$ has two questions and $\mathrm{poly}(1/\varepsilon)$ answers per player; moreover Lemma 4 shows that $\omega^*\big(G_\infty^\delta(\ell)\big) \geq 1 - \varepsilon$. Now, as a direct consequence of Theorem 5 we obtain the following corollary.

Corollary 8. *Let G be a two-player game with two questions per player and finite answer sets $\mathcal{O}_A$ and $\mathcal{O}_B$. Also let $\varepsilon > 0$, $0 < \delta < \omega^*(G) - \omega(G)$ and $\ell = \Theta(\log 1/\varepsilon)$. Any strategy that succeeds at $G_\infty^\delta(\ell)$ with probability at least $\omega^*\big(G_\infty^\delta(\ell)\big) - \varepsilon$ must use a state of local dimension at least $\Omega(1/\sqrt{\varepsilon})$.*

5 Discussion and Open Problems

By exhibiting Hardy's game we have shown that there exist two-party distributed tasks with classical questions and answers for which no finite amount of entanglement allows the parties to perform optimally. We have also exhibited an infinite sequence of strategies, using entanglement of increasing dimension, which obtain success probabilities that tend to the optimal $\omega^*(H) = 1$. Each strategy in our sequence produces answers of fixed length, increasing with the dimension of the entangled state. We have showed that to some extent this is necessary: strategies producing answers with bounded length succeed with probabilities that are bounded away from one. Finally, we have shown that games similar to Hardy's can be constructed out of any game G with two questions per player, finite answer sets and $\omega(G) < \omega^*(G)$.

The most pressing question left open by our work is whether a finite-dimensional shared state is always sufficient to achieve $\omega^*(G)$ for games with finite question and answer sets. As already mentioned in the introduction, there are numerical results which suggest that this is not always the case [PV10]. It is also interesting to investigate what kind of trade-offs may be required in terms of the dimension of strategies which approach the optimum. In particular, are there any general upper bounds on the dimension d sufficient to achieve value $\omega^*(G) - \varepsilon$?

Acknowledgments. The first author thanks the anonymous referees of QIP 2014 for pointing out a mistake in an earlier proof of the non-quantitative part of Theorem 5, Aleksandrs Belovs for suggesting the proof of Lemma 2 as well Richard Cleve, David Roberson, Stephanie Wehner, Harry Buhrman and his group at CWI for helpful discussions. Both authors are grateful to Oded Regev who prompted us to generalize our initial results, which were restricted to Hardy's game, by suggesting that the same construction might also apply for the CHSH game. Both authors are supported by the Ministry of Education, Singapore under the Tier 3 grant MOE2012-T3-1-009. Thomas Vidick is also supported by the Newton Institute, Cambridge. Parts of this work was completed during a workshop at the Institute for Mathematical Sciences, Singapore.

References

[BBT11] Jop Briët, Harry Buhrman, and Ben Toner. A generalized Grothendieck inequality and nonlocal correlations that require high entanglement. *Comm. Math. Phys.*, 305(3):827–843, 2011. arXiv:0901.2009.

[BCMdW10] Harry Buhrman, Richard Cleve, Serge Massar, and Ronald de Wolf. Nonlocality and communication complexity. *Rev. Mod. Phys.*, 82:665–698, 2010. arXiv:0907.3584.

[BCP+14] Nicolas Brunner, Daniel Cavalcanti, Stefano Pironio, Valerio Scarani, and Stephanie Wehner. Bell nonlocality. *Rev. Mod. Phys.*, 86:419–478, 2014. arXiv:1303.2849.

[BPA+08] Nicolas Brunner, Stefano Pironio, Antonio Acin, Nicolas Gisin, André Allan Méthot, and Valerio Scarani. Testing the dimension of Hilbert spaces. *Phys. Rev. Lett.*, 100:210503, May 2008. arXiv:0802.0760.

[BW92] Charles H. Bennett and Stephen J. Wiesner. Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states. *Phys. Rev. Lett.*, 69:2881–2884, 1992. doi:10.1103/PhysRevLett.69.2881.

[BYJK08] Ziv Bar-Yossef, Thathachar S. Jayram, and Iordanis Kerenidis. Exponential separation of quantum and classical one-way communication complexity. *SIAM J. Comput.*, 38(1):366–384, 2008. doi:10.1137/060651835.

[CHTW04] Richard Cleve, Peter Hoyer, Benjamin Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *Proc. of CCC'04*, pages 236–249, 2004. arXiv:quant-ph/0404076.

[Col06] Roger Colbeck. *Quantum And Relativistic Protocols For Secure Multi-Party Computation*. PhD thesis, Trinity College, University of Cambridge, 2006. arXiv:0911.3814.

[CY14] Matthew Coudron and Henry Yuen. Infinite randomness expansion and amplification with a constant number of devices. In *Proc. of STOC'14*, pages 427–436, 2014. arXiv:1310.6755.

[Eke91] Artur K. Ekert. Quantum cryptography based on Bell's theorem. *Phys. Rev. Lett.*, 67(6):661–663, 1991. doi:10.1103/physrevlett.67.661.

[EPR35] Albert Einstein, Boris Podolsky, and Nathan Rosen. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.*, 47:777–780, 1935.

[Fri12] Tobias Fritz. Tsirelson's problem and Kirchberg's conjecture. *Rev. Math. Phys.*, 24(05):1250012, 2012. arXiv:1008.1168.

[Gle57] Andrew M. Gleason. Measures on the closed subspaces of a Hilbert space. *J. Math. Mech.*, 6:885–893, 1957.

[Gol94] Sheldon Goldstein. Nonlocality without inequalities for almost all entangled states for two particles. *Phys. Rev. Lett.*, 72:1951–1951, 1994. doi:10.1103/PhysRevLett.72.1951.

[Har92] Lucien Hardy. Quantum mechanics, local realistic theories, and Lorentz-invariant realistic theories. *Phys. Rev. Lett.*, 68:2981–2984, 1992. doi:10.1103/PhysRevLett.68.2981.

[Har93] Lucien Hardy. Nonlocality for two particles without inequalities for almost all entangled states. *Phys. Rev. Lett.*, 71:1665–1668, 1993. doi:10.1103/PhysRevLett.71.1665.

[IKM09] Tsuyoshi Ito, Hirotada Kobayashi, and Keiji Matsumoto. Oracularization and two-prover one-round interactive proofs against nonlocal strategies. In *Proc. of CCC'09*, pages 217–228, 2009. arXiv:0810.0693.

[KS67] Simon Kochen and Ernst P. Specker. The problem of hidden variables in quantum mechanics. *J. Math. Mech.*, 17:59–87, 1967.

[Lal01] Franck Laloe. Do we really understand quantum mechanics? Strange correlations, paradoxes, and theorems. *Am. J. Phys.*, 69:655, 2001. arXiv:quant-ph/0209123.

[Lot02] M. Lothaire. *Algebraic Combinatorics on Words*. Cambridge University Press, 2002.

[LTW13] Debbie Leung, Ben Toner, and John Watrous. Coherent state exchange in multi-prover quantum interactive proof systems. *Chic. J. Theor. Comput.*, 2013(11):1–18, 2013. arXiv:0804.4118.

[NPA08] Miguel Navascués, Stefano Pironio, and Antonio Acín. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New J. Phys.*, 10(7):073013, 2008. arXiv:0803.4290.

[PAM+10] Stefano Pironio, Antonio Acín, Serge Massar, A. Boyer de La Giroday, Dzimitry N. Matsukevich, Peter Maunz, Steven Olmschenk, David Hayes, Le Luo, T. Andrew Manning, and Cristopher Monroe. Random numbers certified by Bell's theorem. *Nature*, 464(7291):10, 2010. arXiv:0911.3427.

[PV10] Károly F. Pál and Tamás Vértesi. Maximal violation of a bipartite three-setting, two-outcome Bell inequality using infinite-dimensional quantum systems. *Phys. Rev. A*, 82:022116, 2010. arXiv:1006.3032.

[RV13] Oded Regev and Thomas Vidick. Quantum xor games. In *Proc. of CCC'13*, pages 144–155, 2013. arXiv:1207.4939.

[Slo11] William Slofstra. Lower bounds on the entanglement needed to play XOR non-local games. *J. Math. Phys.*, 52(10), 2011. arXiv:1007.2248.

[VV12] Umesh Vazirani and Thomas Vidick. Certifiable quantum dice: or, true random number generation secure against quantum adversaries. In *Proc. of STOC'12*, pages 61–76, 2012. arXiv:1111.6054.

[WCD08] Stephanie Wehner, Matthias Christandl, and Andrew C. Doherty. Lower bound on the dimension of a quantum system given measured data. *Phys. Rev. A*, 78:062112, 2008. arXiv:0808.3960.

[Wil13] Mark M. Wilde. *Quantum Information Theory*. Cambridge University Press, 2013. arXiv:1106.1445.